Application Serial No.	:	09/607,300
Applicants	:	Bret S. Weber et al.
Filed	:	30 June 2000
Title	:	Methods and Structure For An Extensible RAID Storage Architecture
Art Unit	:	2187
Examiner	:	Reba I. Elmore
Docket Number	:	98-063
Date	:	8 June 2004

Mail Stop Appeal Brief - Patents
Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

APPEAL BRIEF

Sir:

Appellants herewith file a Brief in support of their Appeal in the above identified matter.

(1) Real Party in Interest

The real party in interest is LSI Logic, Inc., the employer of the inventors at the time of the invention and the assignee of the patent rights in the above-identified matter.

(2) Related Appeals and Interferences

No other appeals or interferences are known to the appellants, the appellants' legal representative, or the assignee, which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

(3) Status of the Claims

Claims 1-10 and 12-23 (all claims) stand rejected and remain in the application for consideration on appeal.

In the first Office Action mailed 24 September 2002, Examiner Kathy K. Takeguchi objected to various issues in the specification. Examiner Takeguchi also rejected claims 1-6 under 35 U.S.C. §102(e) as anticipated by Walton (United States Patent Number 6,389,494) and rejected claims 7-23 under 35 U.S.C. §103(a) as unpatentable over Walton in various combinations with other references. Appellants

responded to the Office Action on 24 December 2002 addressing the various objections to the specification, discussing the Examiner's rejections and amending independent claims 1, 14 and 19 (and canceling claim 11) to clarify the independence of scalability of the front-end control elements, back-end control elements, and the interconnections therebetween. The Examiner maintained essentially the same rejections (changing the §102(e) rejection to a §103(a) rejection) for essentially the same reasons in a Final Office Action mailed 11 March 2003. The appellants responded to the Final Office Action with further amendments to claims 1-3, 5, 6, 14 and 19 to even more clearly distinguish the cited references. In particular, appellants amended the independent claims substantially to make clear that the interconnect element is a SAN fabric over which messages are exchanged in accord with address indicia within the messages. The amendments also further clarified that the front-end control elements were substantially devoid of circuits for controlling I/O devices (e.g., storage devices such as disk drives) and that the back-end control elements were substantially devoid of circuits for interfacing with host systems (e.g., SCSI or Fibre Channel host interfaces). Examiner Takeguchi responded with an Advisory Action mailed 30 April 2003 indicating a new search would be required for the claims as amended after the final rejection. Appellants filed a Request for Continued Examination mailed 15 May 2003 requesting entry of the prior amendment after final rejection. Examiner Takeguchi then responded with a new, non-final Office Action mailed 14 July 2003 raising new objections to certain claims, rejecting various claims under 35 U.S.C. §112 for use of the term "substantially" and raised the new rejections of all claims under 35 U.S.C. §103(a) as unpatentable over Hashemi (United States Patent Number 5,396,595) in view of Liebhart ("*A Study of an SCI Switch Fabric*", IEEE Modeling, Analysis and Simulation of Computer and Telecommunications Systems 1997, pp. 162-169) and further in view of various other references in combination with Hashemi and Liebhart. Appellants responded on 17 September 2003 discussing Examiner Takeguchi's objections and §112 rejections and amending various claims in response thereto. Appellants also thoroughly discussed the Examiner's new grounds for rejection under §103 raising questions of the Examiner's basis for the rejection and lack of a prima facie case for obviousness. A new examiner, Ms. Reba Elmore, repeated and maintained the §103(a) rejections in a Final Office Action mailed 18 February 2004. The objections

and §112 rejection were withdrawn by Examiner Elmore's Final Office Action. This appeal resulted therefrom.

The 35 U.S.C. §103(a) rejections of all remaining claims (1-10 and 12-23) form the basis of this appeal.

(4) Status of the Amendments

No further amendments have been filed, subsequent to that of 17 September 2003 which responded to the non-final Office Action of 14 July 2003.

(5) Summary of Invention

The invention generally provides a storage subsystem architecture that divides the storage controller function between front-end controllers and back-end controllers and that applies storage area network ("SAN") techniques and devices within the storage subsystem to interconnect the front-end controllers and back-end controllers. The front-end control elements generally perform host interfacing functions while the back-end control elements perform disk I/O processing. SAN components are known and applied outside the storage subsystem for interconnection of such storage subsystems to host computers and other computing subsystems. In the context of this invention, SAN switches are applied ***within*** the storage subsystem to permit more flexible configuration of front-end and back-end control devices within the storage subsystem. By so separating front-end and back-end control functions and coupling the elements with a SAN fabric, improved flexibility and scalability is achieved as compared to the prior art. Each element may be scaled independent of the other elements to better tune performance of the storage system.

A plurality of back-end storage controllers and a plurality of front-end storage controllers are configured within a storage subsystem interconnected by a SAN switching network that permits broad flexibility in interconnecting the various controllers. The front-end controllers ("FECs") are dedicated to "front-end" interfacing to host computer systems and are devoid of circuits and functions to control the disk array devices. The back-end controllers ("BECs") are dedicated to "back-end" control of the disk arrays and

are devoid of circuits and functions to interface directly with the attached host systems. In this architecture, the FECs and BECs are simpler than prior integral controllers that provided both front-end and back-end control functions.

Each FEC and BEC includes a SAN interface to connect to the SAN switched fabric (i.e., to ports of one or more SAN switches). The SAN switches therefore provide flexible interconnection between virtually any number of front-end controllers and any number of back-end controllers. Such a storage subsystem may thereby be flexibly configured to add additional back-end control where required for back-end performance or reliability enhancement independent of the front-end controller configuration and may be configured to add additional front-end controller when required for front-end performance and reliability independent of the back-end controller configuration. Further, the internal SAN fabric may itself be scaled independent of the FEC and BEC elements to provide additional internal communication bandwidth and redundancy.

By providing such configuration flexibility and simpler FEC and BEC devices that segregate their respective functions, the storage subsystem is more scalable than prior known architectures. Additional FECs may be added to alleviate host communication bottlenecks independent of BEC control functions. Conversely, BECs may be added to alleviate disk communication bottlenecks independent of FEC control functions. SAN fabric interconnecting element may be added independent of the FECs and BECs to improve internal communication bandwidth .

In one aspect of the invention, a system is provided with any number of FECs, any number of BECs and a SAN fabric interconnect elements between the FECs and the BECs. Claim 1 broadly recites this aspect as follows:

1. A storage system comprising:

a plurality of front-end control elements devoid of circuits and functions that control a plurality of I/O devices and configured for controlling information exchange using RAID storage management with one or more attached host computer systems;

a plurality of back-end control elements devoid of circuits and functions that interface directly with the attached host computer systems, communicatively coupled to a plurality of I/O devices and configured for controlling information exchange with the I/O devices, wherein the front-end control elements differ in number from the back-end control elements; and

an interconnect element coupled to said front-end control elements and coupled to said back-end control elements to enable exchange of information therebetween, wherein the storage system is adapted to implement additional front-end control elements, back-end control elements and interconnect elements independent of all other such elements wherein the interconnect element is configured for conveying the requests from the front-end control elements to the back-end control elements to perform the host requested I/O operation, and wherein said interconnect element is an SAN architecture fabric that conveys the requests from the front-end control elements to the back-end control elements by exchange of messages between the front-end and back-end control elements according to address indicia within the messages and associated with the front-end and back-end control elements.

Figure 2 shows an exemplary embodiment of these features as storage system 2 that includes a plurality of front-end control elements 200, 208 and 216. Each front-end controller includes one or more front-end interface elements (202, 210 and 218, respectively) to connect the front-end control element (FEC) to one or more host systems 170..174 via a host communication media 160..161. FECs may be connected to a plurality of host system communication media as required for flexible connectivity to host systems.

Each FEC (200, 208, 216) also includes one or more intra-subsystem SAN interfaces (204 and 206, 212 and 214, and 220 and 222, respectively). Intra-subsystem SAN interfaces 204, 206, 212, 214, 220 and 222 are referred to in the subject patent application as "intra-subsystem" to distinguish from SAN interfaces that may be used in a storage subsystem to connect to SAN components external to the storage subsystem. Such external SAN interfaces are not relevant to the operation and structure of the present

invention. As used in the subject patent application "SAN interface" refers to intra-subsystem SAN interfaces as distinct from any SAN interfaces that may be present on a controller for interfacing to host systems external to the storage subsystem.

Each FEC includes one or more SAN interface elements connecting the FEC to the SAN switches 250 and 252 via SAN communication media 254. There are preferably at least two SAN switches 250 and 252 to permit redundant connectivity from the front-end control elements to the plurality of back-end control elements discussed below. There may be any number of such redundant links but in the preferred embodiment, two links from each front-end control element, one to each of two SAN switches, is considered necessary and sufficient. Where reliability of the front-end control communication with the back-end control elements is deemed less important, a single connection between a front-end control element and the SAN switches may be adequate.

Storage system 2 also includes a plurality of back-end control elements 260, 264, 268 and 272 preferably configured as shown in redundant pairs (260 and 264 as a first pair and 268 and 272 as a second pair). Each back-end control element (BEC) includes a SAN interface element (262, 266, 270 and 274, respectively). Each BEC of a redundant pair is connected to one of the two redundant SAN switches 250 and 252 via SAN communication media 256. Specifically as exemplified in figure 2, back-end control element 260 (BEC) connects to SAN switch 250 via SAN interface 262. BEC 264 connects to SAN switch 252 via SAN interface 266. BEC 268 connects to SAN switch 252 via SAN interface 270 and lastly, BEC 272 connects to SAN switch 250 via SAN interface 274.

Each BEC connects to a storage module 280 or 290 comprised of a plurality of disk drives 282 and 292, respectively. Each BEC of a redundant pair preferably connects to one of the storage modules. For example, as shown in figure 2, BEC 260 connects to storage module 280 via media 150 and BEC 264, the other BEC of the redundant pair of 260 and 264, also connects to storage module 280 via media 150. It is also possible for each BEC to provide a pair of redundant links to its associated storage module. For example, as shown in figure 2, redundant pair of BECs 268 and 272 are each redundantly

connected to storage module 290 via a redundant pair of communication links in media 150.

The subject patent application specification expresses the preference for at least pairs of SAN switches and pairs of BECs to ensure redundancy throughout the connections from a host system through to the individual disk drive devices. Any number of FECs, BECs and SAN switches, paired or not, may be configured within the intended scope of the present inventions.

As noted in the subject patent application specification, the SAN switches (250 and 252 of figure 2) and associated SAN communication media 254 and 256 may apply any of several existing SAN architectures. The SAN switches and associated communication media may be any that allows the passing of data and I/O requests between the FECs and the BECs with low latency (i.e., less than 10 microseconds). Typical of such devices/media are local area network (LAN) connections (i.e., Ethernet or Gigabit Ethernet, etc.), Fibre Channel SAN switch devices and media, InfiniBand (see, e.g., www.infinibandta.org) and ServerNet (developed by Tandem and presently sold by HP/Compaq). Though bus structures (such as PCI bus structures) may also be used, the ideal configuration involves a switch that allows for bandwidths that scale with the number of devices (FECs and BECs) that are added to the SAN. By contrast, bus structures are not in general as easily scaled as switched communication media and protocols. Present market forces and technology factors suggest that InfiniBand is a preferred embodiment of the SAN communication media.

In another aspect of the claimed invention, an FEC is provided independent of the BEC and SAN fabric interconnection elements and is configured for coupling through a SAN fabric to a BEC. Claim 14 broadly recites this aspect as follows:

14. A front-end control element for a storage subsystem comprising:

a host system interface;

a processor coupled to said host system interface to process host system generated I/O requests received through said host system interface; and

an SAN interface coupled to said processor for coupling said front-end control element to a plurality of back-end control elements, wherein said front-end control element is adapted to be added to the storage subsystem independent of said back-end control elements,

wherein front-end control elements differ in number from said back-end control elements, and

wherein said SAN interface couples the front-end control element to an SAN fabric that conveys the I/O requests from the front-end control elements to the back-end control elements by exchange of messages between the front-end and back-end control elements according to address indicia within the messages and associated with the front-end and back-end control elements, and

wherein the front-end control element is devoid of circuits and functions that control a plurality of I/O devices.

In another aspect of the claimed invention, a BEC is provided independent of the FEC and SAN fabric interconnect elements and is configured for coupling through a SAN fabric to an FEC. Claim 19 broadly recites this aspect as follows:

19. A back-end control element for a storage subsystem comprising:

a disk drive interface for coupling said back-end control element to a plurality of disk drives; and

an SAN interface coupled to said disk drive interface for coupling said back-end control element to a plurality of front-end control elements, wherein said back-end control element is adapted to be added to the storage subsystem independent of said front-end control elements,

wherein back-end control elements differ in number from said front-end control elements, and

wherein said SAN interface couples the back-end control element to an SAN fabric that conveys the I/O requests from the front-end control elements to the back-end control element by exchange of messages between the front-end and back-end control elements according to address indicia within the messages and associated with the front-end and back-end control elements, and

wherein the back-end control element is devoid of circuits and functions that interface directly with the attached host computer systems.

(6) Issues

Whether claims 1-10 and 12-23 are patentable over the prior art of record.

(7) Grouping of Claims

Group 1: Claims 1-10 and 12-23 stand or fall together sharing the common essential element of independently scalable front-end control and back-end control for a storage subsystem with independently scalable SAN fabric interconnect elements connecting the front and back end control elements.

(8) Argument

As defined in M.P.E.P. §1206, subparagraphs (i), (ii), (iii) and (v), which relate to 35 U.S.C. §112, first paragraph, §112, second paragraph, and §102 and "other common grounds" for rejection, are not relevant to the instant appeal and are therefore omitted. Subparagraph (iv) below discusses the Examiner's remaining rejections of all claims (1-10 and 12-23) under 35 U.S.C. §103(a).

(iv) Rejection of all claims under 35 U.S.C. §103(a)

In the Final Office Action mailed 18 February 2004, the Examiner repeated and maintained the rejection of all claims (1-10 and 12-23) as being unpatentable under 35 U.S.C. §103(a) variously over Hashemi in view of Liebhart and in combination with various other references. In particular:

- claims 1-3, 5-6, 9-10, 13-14, 17-19 and 22-23 were rejected as unpatentable over Hashemi in view of Liebhart,
- claim 4 was rejected over Hashemi in view of Liebhart and further in view of Otteson (United States Patent Number 6,571,310),
- claims 8, 16 and 21 were rejected over Hashemi in view of Liebhart and further in view of Brown (United States Patent Number 6,148,414),

- claims 7, 15 and 20 were rejected over Hashemi in view of Liebhart and further in view of Hennessy (*Computer Architecture: A quantitative Approach;* Morgan Kaufmann Publishers, Inc., Second Edition; 1996; p. 573), and

- claim 12 was rejected over Hashemi in view of Liebhart and further in view of Belsan (United States Patent Number 5,394,532).

At the core of the Examiner's position is the suggestion that Liebhart, when combined with Hashemi, provides the teachings of the recited "SAN fabric" in all claims (i.e., in each independent claim 1, 14 and 19 rejected on the combination of Hashemi with Liebhart). The Examiner properly cites Hashemi as showing a CIM (channel interface module) adapted to interface with host systems and being substantially devoid of circuits to interface with I/O devices (e.g., the rough analog of the recited FEC). The Examiner also properly asserts that Hashemi shows a DIM (device interface module) adapted to interface with I/O devices and substantially devoid of circuits to interface with host systems (e.g., the rough analog of the recited BEC). The CIM and DIM elements each include a multiported buffer memory and processors within each DIM and CIM may share access to the buffer memory of any CIM or DIM. The shared memory access is provided through a FutureBus interconnection.

The Examiner correctly notes that Hashemi fails to show the recited SAN fabric as an interconnect element between the FECs and BECs. Appellants totally agree with the Examiner's observation. FutureBus, or any other bus structure for shared memory access and/or for cache coherency information exchange, does not teach or reasonably suggest use of a storage area network fabric ("SAN fabric") as recited. The Examiner suggests that Liebhart provides such a teaching. Assuming, *arguendo*, that the two references are properly combined by some suggestion found in the art, Liebhart simply fails to teach or reasonably suggest the claimed "SAN fabric" as the Examiner suggests. Rather, Liebhart teaches use of another bus structure for extending shared memory access and/or cache coherency information exchange among a plurality of processors. Since neither Hashemi nor Liebhart teaches or reasonably suggests use of a SAN fabric as an interconnect

element between FECs and BECs of a storage system, neither does the combination teach or reasonably suggest the claimed features.

Specifically, the Examiner points to the Abstract of Liebhart as supporting such a teaching or suggestion to use a SAN fabric as an interconnect element. Neither the Abstract of Liebhart nor the rest of the reference makes any mention of a storage area network fabric (i.e., a "SAN fabric"). Admittedly, the Liebhart reference discusses an SCI bus structure that may be used as a fabric for bus interconnection of computing nodes for purposes of shared memory access and/or cache coherency information exchange. The Abstract of Liebhart is reproduced below in its entirety:

> The Scalable Coherent Interface (SCI) defines a high-speed interconnect system that provides a coherent memory system. It specifies a topology-independent communication protocol with the possibility of connecting up to 64 K nodes. SCI switches are the key components in building large SCI systems effectively. An SCI switch which uses several internal buses is studied, as well as more complex systems composed of several switches. Computer simulations are used to compare the different models and to determine system parameters.

Though appellants can only speculate, it appears that the Examiner relies on the phrase of the Liebhart Abstract relating to topology independence as somehow supporting the purported teaching or suggestion therein of a "SAN fabric." For example, in reference to claim 1 (as well as in rejecting claims 14 and 19), the Examiner states: "The SCI switch fabric is a high speed interconnect system, designed and specified as a topology independent protocol (i.e., interfacable to PCI, VME, Futurebus, ATM, Fibre Channel, etc.)." The reference to "topology" independence seems to suggest that the Examiner believes any medium and any protocol may be adapted for use with the SCI structures and the teachings of Liebhart – including, for example, storage area network protocols and media. Such a suggestion appears to be mere speculation by the Examiner and would not occur to those of ordinary skill in the art. First and foremost, neither reference teaches nor reasonably suggests use of a storage area network ***within*** the storage system as recited in all rejected claims. Liebhart discusses simulated performance testing for a hypothetical SCI switch fabric design. The "topology" of a communication medium has to do with the geometry of interconnections of the various communicating nodes. "Topology" independence does not somehow suggest to one of ordinary skill in

the art that any communication protocol and medium may be equivalently applied nor that the SCI extended bus structure may be applied as a SAN fabric. The Examiner simply labels the SCI switch fabric discussed by Liebhart as a "SAN fabric." Nothing in Liebhart suggests the use of such a processor bus switched fabric for application as a generalized network let alone as a storage area network (SAN).

To the contrary, SCI is clearly distinguished from a network architecture (including a storage area network architecture). The SCI standards relate to processor bus interconnections for purposes of shared memory and cache coherency operations among a plurality of processors. The IEEE standard 1596-1992 is the specification of the SCI (Scalable Coherent Interface). §1.2.1 thereof makes clear that SCI is not a networking standard at all. In pertinent part, the SCI specification statement of scope reads (emphasis added:

> Scope: This standard will encompass two levels of interface, defining operation over distances less than 10 m. The physical layer will specify electrical, mechanical, and thermal characteristics of connectors and cards. The logical level will describe the address space, data transfer protocols, cache coherence mechanisms, synchronization primitives, control and status registers, and initialization and error recovery facilities.
>
> The preceding statements were those submitted to and approved by the IEEE Standards Board as the definition of the SCI project. These goals have been met and exceeded: support for message-passing was added, and the operating distance is not limited to 10 m. (The intent of that limitation was to make clear that his not yet-another Local Area Network).
>
> ***The real distinction between SCI and a network has more to do with the memory-access-based model SCI uses and the distributed cache-coherence model.*** The practical operating distance depends more on the throughput and performance needed than on any absolute limit built into the specification – very long links would yield unacceptable performance for many users (but perhaps not all).

See, IEEE Standard 1596-1992, p. 2. As can be seen from the above excerpt of the SCI standards, SCI is clearly distinguished from network architectures, principally, in that it is oriented to a memory address based model and a cache coherency model – features critical to clustered multi-processing but not at all suggestive of the network features as recited in the rejected claims let alone a storage area network as claimed.

Rather, SCI is a bus protocol that enables the interconnection of processor and memory bus structures of heterogeneous vendors. Although the physical layer for SCI transmission may support high speed serial attachments (such as ATM or Fibre Channel), the SCI specifications make clear that SCI ***is not*** another Local Area Network standard – it is not a network at all! Rather, it is a medium and protocol by which processor/memory bus structures may be extended so that shared memories may be used and cache coherency may be maintained among a plurality of processing nodes using (potentially) disparate processor bus structures.

While it may be proper to refer to Liebhart's hypothetical switched structure as a "fabric", it is clearly not a storage area network (SAN) fabric as recited in the rejected claims. One widely used technical dictionary defines a "storage area network" as follows:

> A SAN is a special purpose high-speed network designed to transport database-intensive applications, such as those used for inventory, billing, receivables, customer relationship management and supply chain. The concept is to have a dedicated network for these applications so users get fast response and don't get bogged down in the corporation's general networking management and supply chain. Actually a SAN generally is in the form of a sub-network that is part of a larger LAN (Local Area Network). The network protocols can include ATM (Asynchronous Transfer Mode), ESCON (Enterprise Systems Connectivity), Fast Ethernet (100 Mbps or Gigabit), FC-AL (Fibre Channel-Arbitrated Loop), or SSA (Serial Systems Architecture).The storage can be JBOD (Just a Bunch Of Disks), RAID (Redundant Array of Inexpensive Disks), a bunch of servers on a network, or a more complex and expensive host storage server such as a midrange or mainframe computer. SAN applications include disk mirroring, data backup and restoration, data archival and retrieval, data transfer between storage devices, and data sharing between servers. A SAN is much more complex than simple Network Attached Storage (NAS).

See, *Newton's Telecom Dictionary*; CMP Books; Feb. 2002; p. 645. Another technical glossary definition reads:

> A storage area network (SAN) is a high-speed special-purpose network (or subnetwork) that interconnects different kinds of data storage devices with associated data servers on behalf of a larger network of users. Typically, a storage area network is part of the overall network of computing resources for an enterprise. A storage area network is usually clustered in close proximity to other computing resources such as IBM z990 mainframes but may also extend to

> remote locations for backup and archival storage, using wide area network carrier technologies such as ATM or SONET.
>
> A storage area network can use existing communication technology such as IBM's optical fiber ESCON or it may use the newer Fibre Channel technology. Some SAN system integrators liken it to the common storage bus (flow of data) in a personal computer that is shared by different kinds of storage devices such as a hard disk or a CD-ROM player.
>
> SANs support disk mirroring, backup and restore, archival and retrieval of archived data, data migration from one storage device to another, and the sharing of data among different servers in a network. SANs can incorporate subnetworks with network-attached storage (NAS) systems.
>
> See, the "Glossary" link at http://searchstorage.techtarget.com.

These common definitions of SAN reflect the understanding of those of ordinary skill in the art that a storage area network is first and foremost a network – a structure distinguished from the shared memory/cache coherency bus structures supported by FutureBus or SCI as taught by Hashemi and/or Liebhart. The subject application acknowledges this common understanding of a SAN (and hence a SAN fabric) and provides for use of such a SAN fabric ***within*** a storage system as an interconnect element between front-end and back-end control elements.

Since Hashemi teaches the use of shared buffer memories between the CIM and DIM modules, it may arguably be proper to combine the teachings of Liebhart with that of Hashemi to improve the shared memory features of Hashemi. It may properly be asserted that Hashemi teaches the use of a shared memory bus structure (specifically FutureBus) to permit controllers (CIMs and DIMs) within a storage system to access shared buffer memories on each controller. It is further reasonable to suggest that Hashemi in view of Liebhart teaches the extension of the FutureBus structure to use a hypothetical SCI fabric as discussed in Liebhart. This combination may be proper in that SCI specifications anticipate precisely such an extension of a FutureBus (or other bus architectures) to permit enhanced access to shared memories or to enhance cache coherency processing among a plurality of processors. However, nothing in Hashemi or Liebhart, considered individually or in combination, teaches or reasonably suggests the

use of a network structure between the FEC and BEC elements as recited in the rejected claims.

In sum, Hashemi and Liebhart, considered individually or in combination fail to teach or reasonably suggest provision of a SAN fabric as an interconnect element in a storage system between front-end and back-end control elements as recited in rejected claim 1. Neither do these references, considered individually or in combination, teach or reasonably suggest a front-end control element adapted for coupling to back-end control elements through such a SAN fabric or a back-end control element adapted for coupling to a front-end control element through such a SAN fabric. Thus, for similar reasons to those discussed above, the references, considered individually or in combination, do not teach or reasonably suggest the structures of rejected claims 14 and 19.

All dependent claims were rejected based either upon this same combination of Hashemi and Liebhart or upon this combination further in view of other references. For at least the same reasons as presented above and as dependent from allowable base claims, appellants maintain that claims 2-10, 12-13, 15-18 and 20-23 are allowable over all prior art of record, considered individually or in any combination.

Summary

Appellants argue that the Examiner's rejection of all claims (1-10 and 12-23) under 35 U.S.C. §103(a) is inadequate as a matter of law and should be reversed. It is believed that this Appeal Brief has been timely filed within two (2) months of receipt of the Notice of Appeal (mailed 13 April 2004). However, if an extension of time is deemed to be required by the Patent Office, the Patent Office is hereby requested to accept this request as a petition for a one (1) month extension of time to respond with any requisite fees therefore being charged to deposit account 12-2252.

Respectfully submitted,

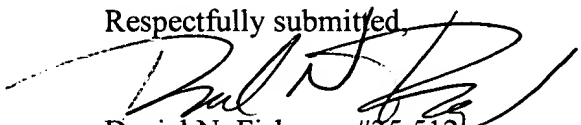

Daniel N. Fishman #35,512
Duft Setter Ollila & Bornsen
2060 Broadway, Suite 300
Boulder, CO 80302
(303) 938-9999 x21
(303) 938-9995 (fax)

(9) Appendix

Current Claims

1. A storage system comprising:

a plurality of front-end control elements devoid of circuits and functions that control a plurality of I/O devices and configured for controlling information exchange using RAID storage management with one or more attached host computer systems;

a plurality of back-end control elements devoid of circuits and functions that interface directly with the attached host computer systems, communicatively coupled to a plurality of I/O devices and configured for controlling information exchange with the I/O devices, wherein the front-end control elements differ in number from the back-end control elements; and

an interconnect element coupled to said front-end control elements and coupled to said back-end control elements to enable exchange of information therebetween, wherein the storage system is adapted to implement additional front-end control elements, back-end control elements and interconnect elements independent of all other such elements wherein the interconnect element is configured for conveying the requests from the front-end control elements to the back-end control elements to perform the host requested I/O operation, and wherein said interconnect element is an SAN architecture fabric that conveys the requests from the front-end control elements to the back-end control elements by exchange of messages between the front-end and back-end control elements according to address indicia within the messages and associated with the front-end and back-end control elements.

2. The system of claim 1 further comprising:

a plurality of disk drives coupled as I/O devices to said back-end control elements.

3. The system of claim 2

wherein said plurality of disk drives comprises:

a first subset of said plurality of disk drives; and

a second subset of said plurality of disk drives, and

wherein said plurality of back-end control elements includes:

a first pair of back-end controllers coupled to said first subset; and

a second pair of back-end controllers coupled to said second subset.

4. The system of claim 3 further comprising:

a first pair of redundant links coupling said first pair of back-end controllers to said first subset; and

a second pair of redundant links coupling said second pair of back-end controllers to said second subset.

5. The system of claim 1 wherein said interconnect element comprises a pair of interconnect elements and wherein each of said plurality of front-end control elements is coupled to each of said pair of interconnect elements.

6. The system of claim 5 further comprising:

a first set of disk drives; and

a second set of disk drives, said plurality of back-end control elements including:

a first pair of back-end controllers coupled to said first set wherein each of said first pair of back-end controllers is coupled to a corresponding one of said pair of interconnect elements; and

a second pair of back-end controllers coupled to said second set wherein each of said second pair of back-end controllers is coupled to a corresponding one of said pair of interconnect elements.

7. The system of claim 1 wherein said interconnect element comprises a PCI bus.

8. The system of claim 1 wherein said interconnect element comprises:

a Fibre Channel communication medium; and

a Fibre Channel SAN switch coupled to said Fibre Channel communication medium.

9. The system of claim 1 wherein said interconnect element comprises an InfiniBand compliant communication medium.

10. The system of claim 1 wherein said interconnect element comprises a local area network communication medium.

11. (Cancelled)

12. The system of claim 1 wherein said front-end control element is operable to perform mapping of logical storage addresses to physical storage addresses for further operations by said back-end control element.

13. The system of claim 1 wherein said back-end control further comprises:
a RAID parity assist element for RAID parity generation and checking.

14. A front-end control element for a storage subsystem comprising:

a host system interface;

a processor coupled to said host system interface to process host system generated I/O requests received through said host system interface; and

an SAN interface coupled to said processor for coupling said front-end control element to a plurality of back-end control elements, wherein said front-end control element is adapted to be added to the storage subsystem independent of said back-end control elements,

wherein front-end control elements differ in number from said back-end control elements, and

wherein said SAN interface couples the front-end control element to an SAN fabric that conveys the I/O requests from the front-end control elements to the back-end control elements by exchange of messages between the front-end and back-end control elements according to address indicia within the messages and associated with the front-end and back-end control elements, and

wherein the front-end control element is devoid of circuits and functions that control a plurality of I/O devices.

15. The front-end control element of claim 14 wherein said SAN interface comprises a PCI bus interface.

16. The front-end control element of claim 14 wherein said SAN interface comprises a Fibre Channel communication media interface.

17. The front-end control element of claim 14 wherein said SAN interface comprises an InfiniBand compliant communication medium.

18. The front-end control element 14 wherein said SAN interface comprises a local area network communication medium.

19. A back-end control element for a storage subsystem comprising:

a disk drive interface for coupling said back-end control element to a plurality of disk drives; and

an SAN interface coupled to said disk drive interface for coupling said back-end control element to a plurality of front-end control elements, wherein said back-end control element is adapted to be added to the storage subsystem independent of said front-end control elements,

wherein back-end control elements differ in number from said front-end control elements, and

wherein said SAN interface couples the back-end control element to an SAN fabric that conveys the I/O requests from the front-end control elements to the back-end control element by exchange of messages between the front-end and back-end control elements according to address indicia within the messages and associated with the front-end and back-end control elements, and

wherein the back-end control element is devoid of circuits and functions that interface directly with the attached host computer systems.

20. The back-end control element of claim 19 wherein said SAN interface comprises a PCI bus interface.

21. The back-end control element of claim 19 wherein said SAN interface comprises a Fibre Channel communication media interface.

22. The back-end control element of claim 19 wherein said SAN interface comprises an InfiniBand compliant communication medium.

23. The back-end control element of claim 19 wherein said SAN interface comprises a local area network communication medium.